

at least a second semantic data context associated with at least one of the heterogeneous data sources. The data translator translates received data from the semantic data contexts associated with the heterogeneous data sources into the semantic data context associated with the request. As the specification points out, an associated semantic data context is the set of assumptions made by a data source or receiver regarding the meaning of data, e.g., the units in which the data is stored. Therefore, independent claim 1 recites a system which can accept a request having a number of assumptions about how data is represented, translate that request into a request that matches the assumptions made by a data source (e.g., a database) about how data is stored, and translate the data from the data source into a form that matches the assumptions of the requester.

Independent claim 14 recites a method of querying heterogeneous data sources over a network which comprises the steps of translating a request having an associated semantic data context into a query having at least a second semantic data context associated with at least one of the heterogeneous data sources to be queried and translating received data from the semantic data contexts associated with the heterogeneous data sources into the semantic data context associated with the request. Independent claim 26 recites an article of manufacture having computer-readable program means for performing similar steps. Thus, each system or method recited by the independent claims provides a mechanism for translating between the meaning assigned to data values by disparate databases.

Rejection of Claims 1-38 Under 35 U.S.C. §103(a)

The Examiner rejected claims 1-38 as obvious over U.S. Patent No. 5,416,917 to Adair et al. ("the '917 patent") alone, or in view of U.S. Patent No. 5,590,319 to Cohen et al. ("the '319 patent"). Applicant respectfully traverses this rejection to the extent it is maintained over the claims as amended.

The '319 patent describes a query processor which breaks an input query into a plurality of parallel output queries that are directed to individual databases or partitions of

those databases in order to take advantage of parallel processing techniques. Although the '319 patent refers to querying heterogeneous databases, it is clear that the '319 patent uses this term to refer to databases that are accessed using different dialects of a single query language or completely different query languages (the '319 patent, Col. 3, Lns. 55-58), that is, the '319 patent refers to databases that are syntactically heterogeneous, not semantically heterogeneous. Once syntactically heterogeneous databases are queried using the apparatus of the '319 patent, a "runner" processes the returned result sets and produces answer sets satisfying the source query. However, the runner performs no translation between the semantic data values returned by the databases and the semantic data values expected by the query. The runner may simply relay output rows more or less unchanged to the client; the runner may merge answer sets returned from different databases; the runner may remove duplicate rows as it merges answer sets; the runner may apply SQL aggregate functions to column values (i.e. sum); the runner may combine rows that satisfy a specified join condition; or the runner may write a result into a designated memory location (the '319 patent, Col. 14, Lns. 7-17).

Applicants respectfully submit that not only does the '319 patent fail to disclose a system for querying heterogeneous data sources which includes semantic data context interchange, it also fails to suggest such a system. For example, the sample query beginning at Col. 9, Ln. 26 of the '319 patent and continuing through Col. 10, Lns. 15 of the '319 patent illustrates that the '319 patent conceives only of databases which have homogenous semantic data context. In this example, a query is normalized into a query that includes a restriction that month must equal "DEC". This statement is included throughout the various normalized subqueries (see Col. 9, Ln. 57; Col. 9, Ln. 65; Col. 10, Ln. 7; and Col. 10, Ln. 14). Accordingly, the '319 patent would fail to retrieve sales data from a semantically different database, such as one in which month names are entered in German, which would abbreviate "December" to "DEZ," not "DEC."

Accordingly, Applicants respectfully submit that the '319 patent neither teaches nor suggests the invention recited by amended independent claims 1, 14, and 26 and that the rejection of claims 1-38 as obvious over the '319 patent alone has been overcome.

The '917 patent describes a database system in which individual machines running SQL having different syntactic representations for data values having SQL data types (the '917 patent, Col. 4, Lns. 39-44). The method and apparatus described by the '917 patent allow a query to be addressed to multiple databases, each of which stores its data using a different bit-level representation from the other machines in the system. Thus, the '917 patent discloses only systems or methods which translate between the form of physical data storage of each database, not the semantic data context as claimed by the instant application. Accordingly, Applicants respectfully submit that the '917 patent neither teaches nor suggests the invention recited by amended independent claims 1, 14, and 26, and that the rejection of claims 1 through 38 as obvious in view of the '917 patent alone has been overcome.

Since neither the '917 patent nor the '319 patent teach or suggest a system or method in which a request having an associated semantic data context is translated into a query having a semantic data context associated with at least one of the data sources, Applicants respectfully submit that the combination of those two references would similarly lack such a feature. Accordingly, Applicants submit that the rejection of claims 1-38 as obvious over the '319 patent and the '917 patent in combination has been overcome.

CONCLUSION

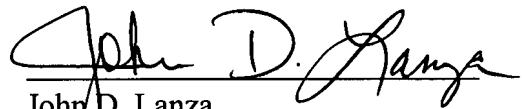
In view of the arguments made above, Applicants respectfully submits that the claims are in condition for allowance and request that the claim rejections be reconsidered and withdrawn.

If the Examiner believes that a telephone conversation with Applicants' attorney would expedite allowance of this application, the Examiner is cordially invited to call the undersigned attorney at (617) 248-7604.

Respectfully submitted,

Date: April 7, 1998
Reg. No. 40,060

Tel. No.: (617) 248-7604
Fax No.: (617) 248-7100
412JDL5473/60.476917-1

A handwritten signature in cursive script, reading "John D. Lanza". The signature is written in dark ink and is positioned above the printed name and address.

John D. Lanza
Attorney for Applicants
TESTA, HURWITZ & THIBEAULT, LLP
High Street Tower
125 High Street
Boston, MA 02110